

# Beyond Hybrid Networking

Cees de Laat

University of Amsterdam



# Trends

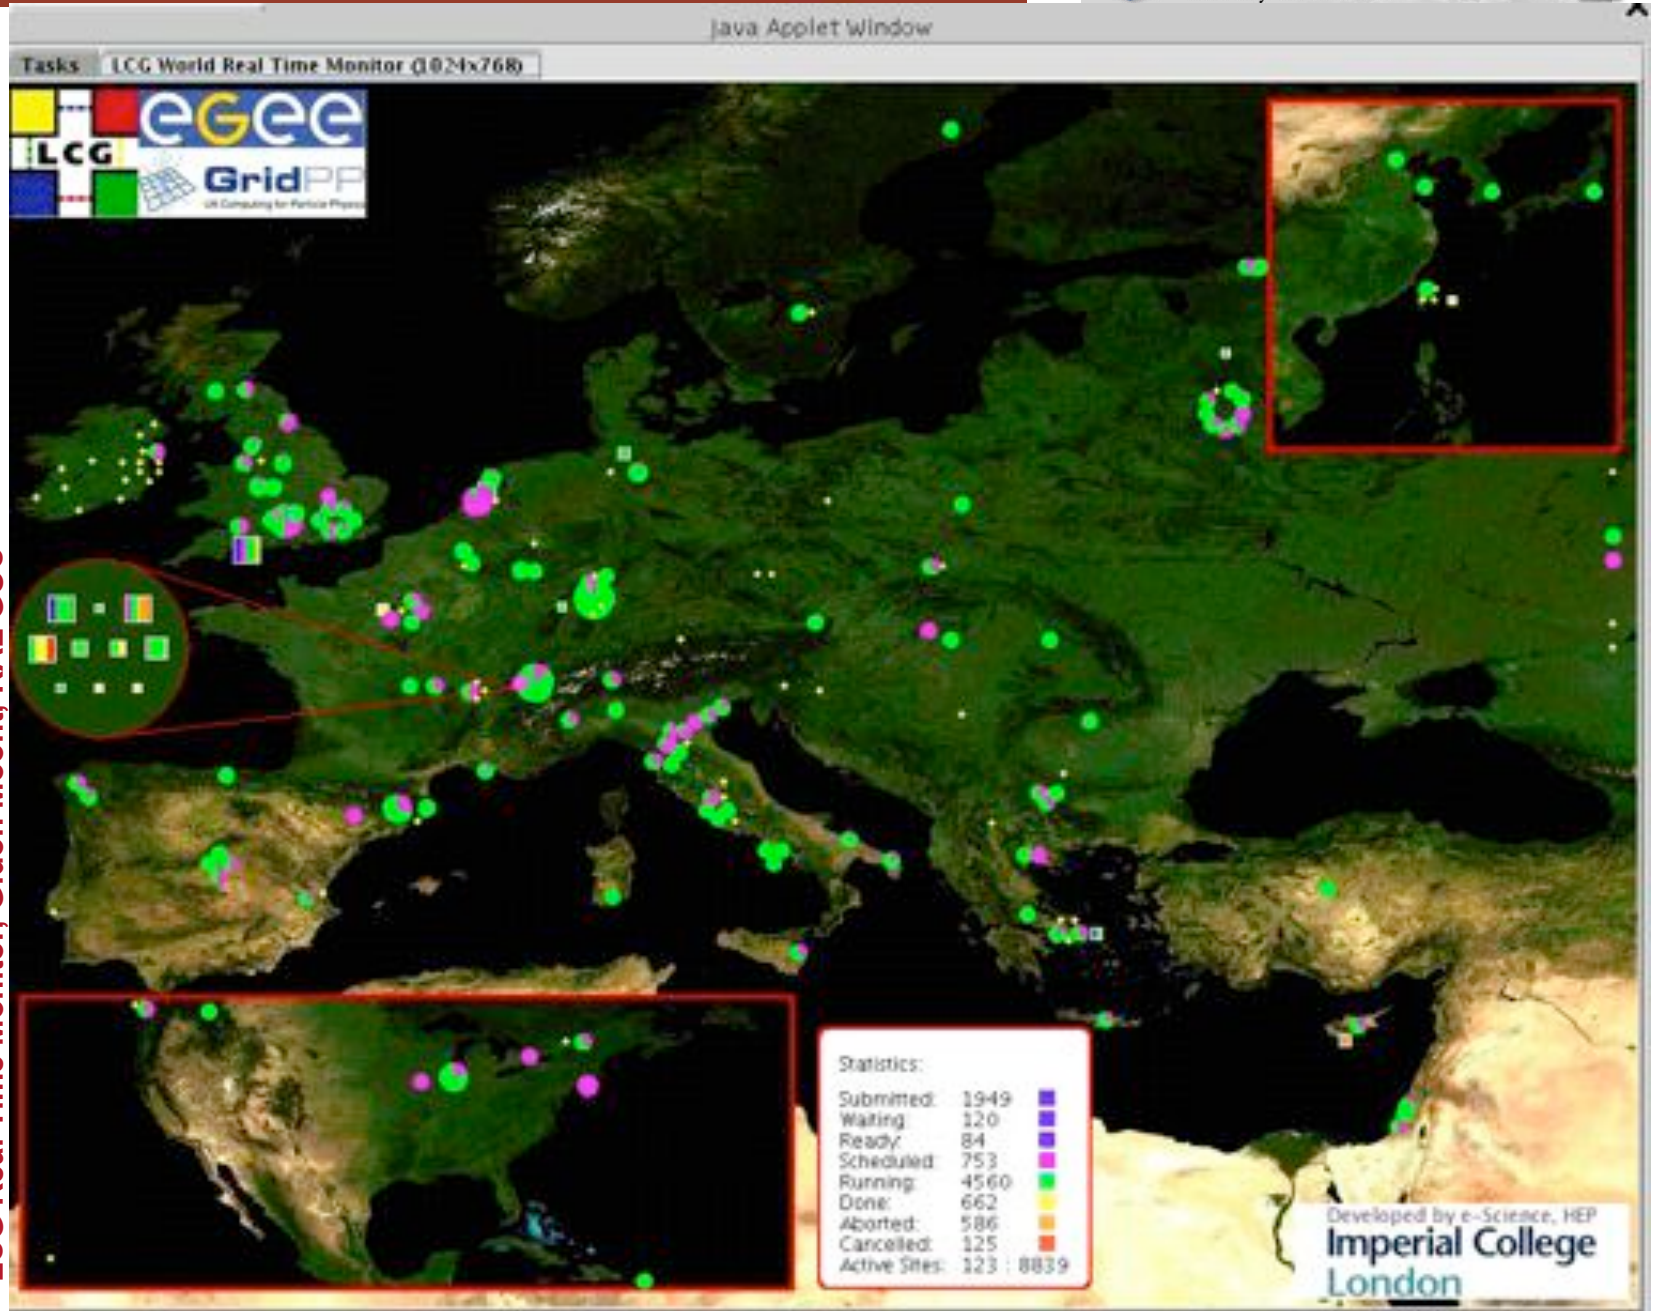
- We have made baby-steps on the path to optical networking
  - Still many mails and phone calls
- See several trends:
  - lambda's get fatter and cheaper
  - photonic technology cheap per bandwidth
  - 100 Gb trials ongoing
  - embedded computation capacity increasing
  - latency and high bandwidth congestion avoidance conflict
  - ethernet is enhanced with circuit properties (PBT)
  - applications need more and more predictable behaviour





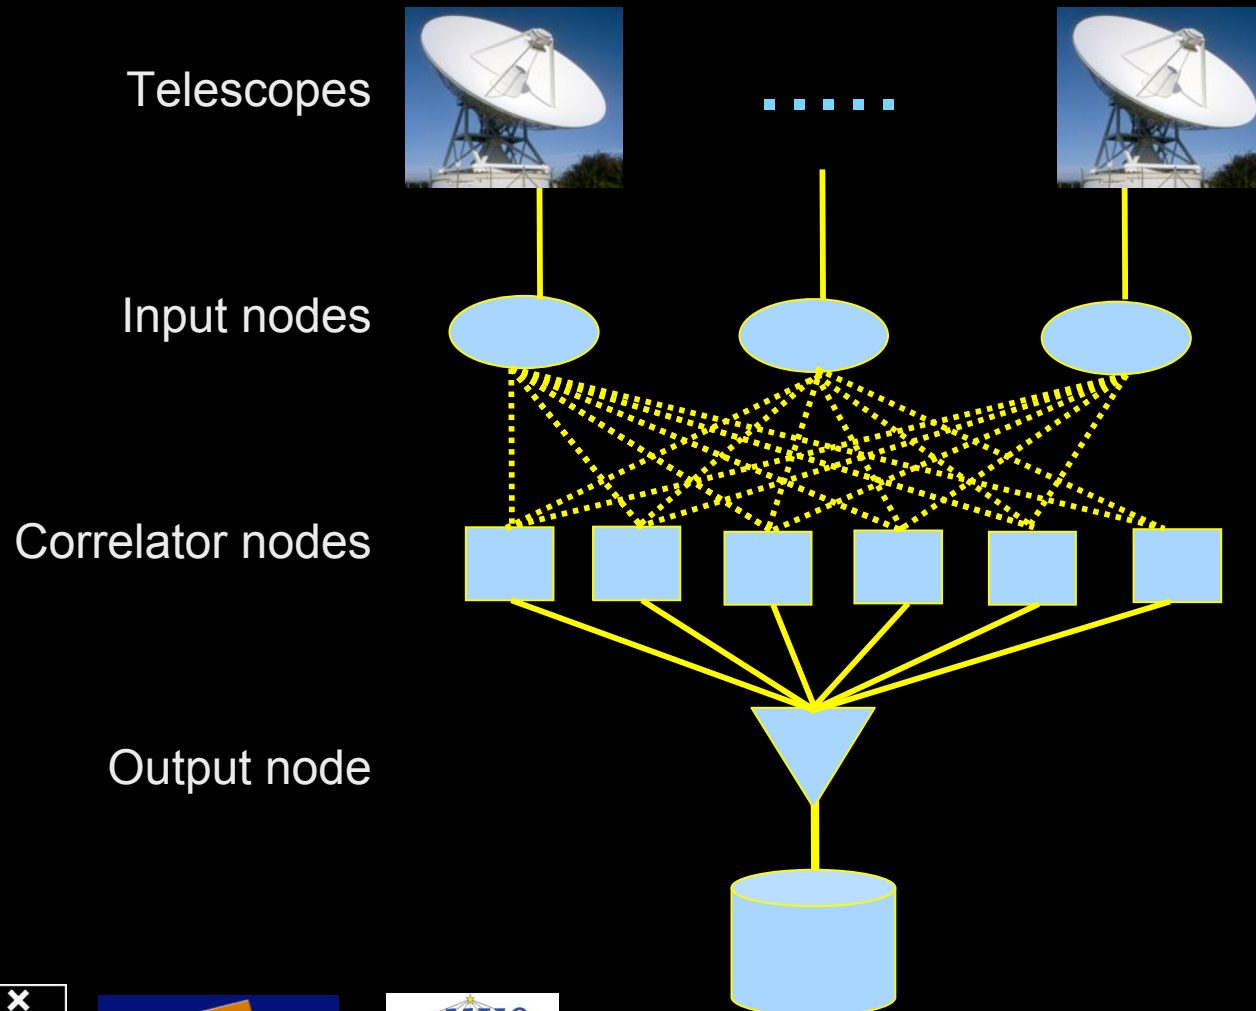
# Grid in operation

LCG Real Time Monitor, Gidon Moont, RAL GOC



# The SCARIE project

**SCARIE:** a research project to create a Software Correlator for e-VLBI.  
**VLBI Correlation:** signal processing technique to get high precision image from spatially distributed radio-telescope.



To equal the hardware correlator we need:

16 streams of 1Gbps

16 \* 1Gbps of data

2 Tflops CPU power

2 TFlop / 16 Gbps =

**1000 flops/byte**

**THIS IS A DATA  
FLOW PROBLEM !!!**





# CineGrid



# CineGrid Format - Numbers - Bits

Format	X	Y	Rate	Color bits/pix	Frame pix	Frame MByte	Flow MByt/s	Stream Gbit/s
720p HD	1280	720	60	24	921600	2.8	170	1.3
1080p HD	1920	1080	30	24	2073600	6.2	190	1.5
2k	2048	1080	24 48	36	2211840	10	240 480	1.2 2.4
SHD	3840	2160	30	24	8294400	25	750	6.0
4k	4096	2160	24	36	8847360	40	960	7.6

Note: this is excluding sound!


Note: these are raw uncompressed data rates!







# Multi-layer Network Description Language





# TeraThinking

- What constitutes a Tb/s network?
- CALIT2 has 8000 Gigabit drops ?->? Terabit Lan?
- look at 80 core Intel processor
  - cut it in two, left and right communicate 8 TB/s
- think back to teraflop computing!
  - MPI makes it a teraflop machine
- massive parallel channels in hosts, NIC's
- TeraApps programming model supported by
  - TFlops -> MPI / Globus
  - TBytes -> OGSA/DAIS
  - TPixels -> SAGE
  - TSensors -> LOFAR, LHC, LOOKING, CineGrid, ...
  - Tbit/s -> ?



# Need for discrete parallelism

- it takes a core to receive 1 or 10 Gbit/s in a computer
- it takes one or two cores to deal with 10 Gbit/s storage
- same for Gigapixels
- same for 100's of Gflops
- Capacity of every part in a system seems of same scale
- look at 80 core Intel processor
  - cut it in two, left and right communicate 8 TB/s
- massive parallel channels in hosts, NIC's
- Therefore we need to go massively parallel allocating complete parts for the problem at hand!





# Mathematica enables advanced graph queries, visualizations and real-time network manipulations on UPVNs

Topology matters can be dealt with algorithmically

Results can be persisted using a transaction service built in UPVN

## Initialization and BFS discovery of NEs

```
Needs["WebServices`"]
<<DiscreteMath`Combinatorica`
<<DiscreteMath`GraphPlot`
InitNetworkTopologyService["edge.ict.tno.nl"]
```

Available methods:

```
{DiscoverNetworkElements, GetLinkBandwidth, GetAllIpLinks, Remote,
NetworkTokenTransaction}
```

Global`upvnverbose = True;

```
AbsoluteTiming[nes = BFSDiscover["139.63.145.94"];][[1]]
```

```
AbsoluteTiming[result = BFSDiscoverLinks["139.63.145.94", nes];][[1]]
```


Getting neighbours of: 139.63.145.94

Internal links: {192.168.0.1, 139.63.145.94}

(...)

Getting neighbours of: 192.168.2.3

Internal links: {192.168.2.3}



## Transaction on shortest path with tokens

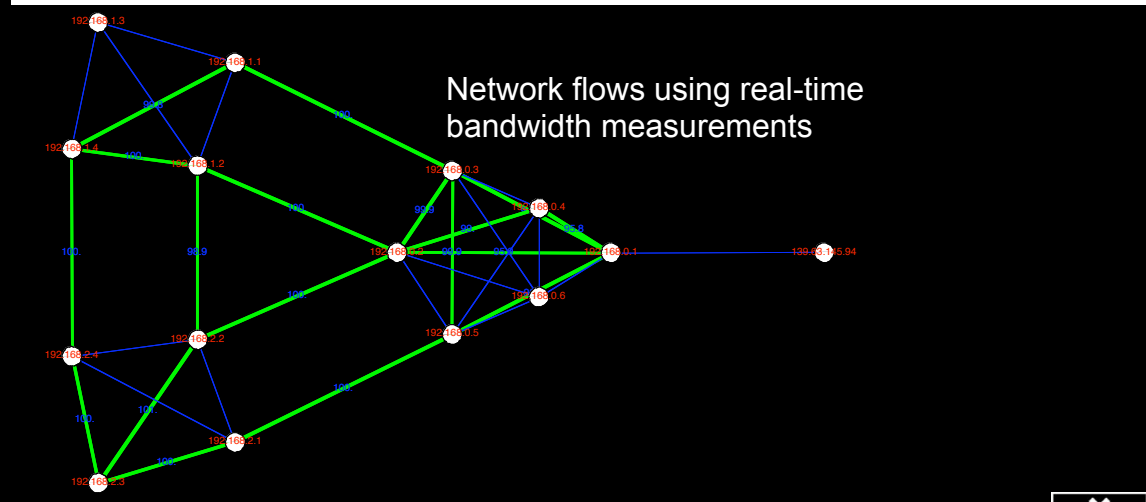
```
nodePath = ConvertIndicesToNodes[
  ShortestPath[ g,
    Node2Index[nids, "192.168.3.4"],
    Node2Index[nids, "139.63.77.49"],
    nids];
```

```
Print["Path: ", nodePath];
If[NetworkTokenTransaction[nodePath, "green"]==True,
  Print["Committed"], Print["Transaction failed!"]];
```

Path:

```
{192.168.3.4, 192.168.3.1, 139.63.77.30, 139.63.77.49}
```

Committed



ref: Robert J. Meijer, Rudolf J. Strijkers, Leon Gommans, Cees de Laat, User Programmable Virtualized Networks, accepted for publication to the IEEE e-Science 2006 conference Amsterdam.


StarPlane







# Multi Layer Service Architecture



# Recommendations

- 1) The current draft fire report lacks vision, goals, does not define stakeholders, defines testbed but for what? Why?
- 2) Let research determine the testbed, not other way around! Allow only testbed if R&D, goals, vision is defined.
- 3) Look at the 10 year timescale in Internet development, underlying methods (paradigms) and technology.
- 4) 100 Gbit/s - Tb/s nets, look at grooming!
- 5) Populate the infrastructure with programmable L1, L2 and L3 objects and the virtualization towards applications
- 6) take into account the new internet exchanges (GOLE's)
- 7) Security infrastructure  $\Leftrightarrow$  infrastructure security
- 8) Do not demand too many different details in a proposal. Writing a proposal is now already a huge effort.



# *Questions ?*

I did not talk about *StarPlane*



**AAA  
Security**

...

